

Measuring Traffic on the Wireless Medium: Experience and Pitfalls

Jihwang Yeo, Suman Banerjee, Ashok Agrawala

Department of Computer Science

University of Maryland

College Park, MD 20742

{jyeo, suman, agrawala}@cs.umd.edu

CS-TR-4421 and UMIACS-TR-2002-101

December 17, 2002

Abstract

A number of measurement studies have examined traffic characteristics in wireless networks. Most of these measurements [1], [6], [7] have been conducted from the wired portion of the network. In this paper we argue that such measurements are not sufficient to expose either the characteristics of the wireless medium or how such characteristics impact traffic patterns. While it is easier to make consistent measurements in the wired part of a network, such measurements can not observe the significant vagaries present in the wireless medium itself. As a consequence constructing an efficient and accurate measurement system from a wireless vantage point is important but usually quite difficult. In our work we have explored the various issues in implementing such a system to monitor traffic in an 802.11 based wireless network. We identify different challenges in making such measurements and provide detailed experimental evidence in their supports.

Our work shows that the wireless measurement allows us to infer much richer information about the medium characteristics than is possible with a measurements made on the wired part of the network. We apply our measurement technique to study the end-to-end wireless network delay. We show that wireless monitoring can effectively identify the causes of end-to-end delays.

I. INTRODUCTION

With the popularity of 802.11 based wireless networks, it becomes increasingly important to understand the characteristics of the wireless traffic and wireless medium itself. Such wireless networks use a wired connection also, which is typically used to connect the access points of the 802.11 network to the Internet, for example. A number of measurement studies [1], [6], [7], [3], [5] have examined traffic characteristics in wireless networks. In these studies, the measurements have been conducted on the wired portion of the network.

Measurements with a wired connection can provide accurate traffic statistics as seen in that portion of the network. However they are mostly unable to expose the wireless medium characteristics because they cannot see the actual frames on the air. In this paper we argue that traffic measurements from a wired vantage point in the network are not sufficient to expose either the wireless medium characteristics or how such characteristics impact the traffic patterns.

In our study we examine the following two questions: Is it possible to make meaningful measurements of the wireless medium using a wireless measurement infrastructure? If this is possible, what are the pitfalls that a wireless measurement system needs to be aware of so that the observations are consistent with the real world wireless experience?

We have performed a detailed study over a period of three months in which we have observed the wireless characteristics in the A.V. Williams building on the campus of the University of Maryland, which houses the

This work was supported in part by the Maryland Information and Network Dynamics (MIND) Laboratory, its Founding Partner Fujitsu Laboratories of America, and by the Department of Defense through a University of Maryland Institute for Advanced Computer Studies (UMIACS) contract.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 17 DEC 2002		2. REPORT TYPE		3. DATES COVERED 00-00-2002 to 00-00-2002	
4. TITLE AND SUBTITLE Measuring Traffic on the Wireless Medium: Experience and Pitfalls				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Maryland, Department of Computer Science, College Park, MD, 20742				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT A number of measurement studies have examined traffic characteristics in wireless networks. Most of these measurements [1], [6], [7] have been conducted from the wired portion of the network. In this paper we argue that such measurements are not sufficient to expose either the characteristics of the wireless medium or how such characteristics impact traffic patterns. While it is easier to make consistent measurements in the wired part of a network, such measurements can not observe the significant vagaries present in the wireless medium itself. As a consequence constructing an efficient and accurate measurement system from a wireless vantage point is important but usually quite difficult. In our work we have explored the various issues in implementing such a system to monitor traffic in an 802.11 based wireless network. We identify different challenges in making such measurements and provide detailed experimental evidence in their supports. Our work shows that the wireless measurement allows us to infer much richer information about the medium characteristics than is possible with a measurements made on the wired part of the network. We apply our measurement technique to study the end-to-end wireless network delay. We show that wireless monitoring can effectively identify the causes of end-to-end delays.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Department of Computer Science. The network there is known to have a high load of wireless traffic. Our observations indicate that indeed such a wireless measurement based inference is possible. However, the measurement process is significantly more challenging than when performed from a wired vantage point. At the same time, we notice that wireless measurements provide much richer information about the medium characteristics than is possible with measurements on the wired connection. We discuss some of applications that can exploit the information.

A. Advantages of Wireless Traffic Monitoring

Wireless monitoring (which we refer to in this paper as, “sniffing”) is useful to understand the traffic characteristics in wireless network for the following reasons.

A sniffing system can be easily set up and put into operation without any interference to existing infrastructure, including end-hosts or network routers. In fact sniffing can be performed without any interaction with the existing network, and hence is completely independent of the operational network.

Another reason is that we can have more control over the geographical coverage of the measurement. The sniffing devices theoretically can listen to any frames on the air within the range. Therefore we can tailor the measurement coverage for our purpose by placing the sniffers over the proper area.

Most importantly, sniffing can provide the rich information on the wireless medium itself so that we can infer the wireless medium characteristics. Such information consists of physical and link layer information of each packet, which are not available if traffic monitoring is performed in the wired part of the network. Thus wireless sniffing allows the measurement system to examine physical layer header information including signal strength, noise level and data rate for individual packets. Similarly the link layer headers carry information which include 802.11 type and control fields [4].

Those physical and link layer data, collected by sniffers, can reveal more than overall traffic statistics in several ways.

First, 802.11 control and management frames, (for example beacons, RTS, CTS and ACKs) provide the local information of each network node as well as the global network status. Beacon frames are used for global clock synchronization by an AP. RTS/CTS and ACKs are exchanged for medium reservation and transfer notification respectively [4].

Second, combined with timestamps the collected data can be used as good traces of 802.11 link-level operations. Such traces are useful when we want to emulate the protocol or diagnose problems of wireless networks.

Third, sniffer traces can be used to detect the misbehaviors of some network nodes. With such traces, we can check how well each node conforms to 802.11 protocol. If any misbehaving nodes are found, based on the traces we can infer the clue to how to handle the problematic nodes.

B. Challenges of Wireless Monitoring

The advantages we mentioned above would not be meaningful, unless our sniffer can capture nearly all the frames on the air. Unfortunately it is very difficult to guarantee that the sniffers can see 100% of wireless frames. In fact it is even difficult to correctly estimate the number of packets different sniffers lose due to differences in various wireless cards, drivers, and antennae characteristics. We go over this problem in Section III.

Losses in the sniffers poses the most challenging problem in wireless monitoring. There are several categories of losses, *frame loss*, *type loss* and *AP loss*. By *frame loss* we mean existence of such frames that are present on the air at the time of a measurement, but are not detected by the sniffers. If we say a sniffer has *type loss*, we mean the sniffer is unable to capture specific types of packets inherently. Similarly *AP loss* of a sniffer means that the sniffer loses nearly all the frames originating from specific APs.

We have observed that typically most of these losses are due to signal strength variability, card variability or a combination of both. For example as signal condition becomes worse, a wireless sniffing device is more susceptible to frame losses. Some specific card implementations do not allow ACK frames to be passed up

to user applications, and hence type loss results. AP loss sometimes occurs due to incompatibility between AP firmware and user card firmware.

In this work, we present the details on the observations of sniffer loss and its loss variability in Section III. We observe various kinds of losses and their variability in the sniffers during our experiments. While most of our experiments perform passive measurements, we present some active experiment results to explain how signal strength and card variability cause packet loss and end-to-end delay variability between the AP and a wireless station in Section IV.

C. Key Contributions

In this work we have extensively studied passive measurements of the wireless medium. We have used a number of different measurement equipments which provide some interesting insights. We have repeated the same experiments with various setups and obtain the observations on sniffer loss and loss variability. Our findings can be summarized as follows:

- Percentage of valid 802.11 frames that are not detected by the sniffing devices (which we call lost frames) can increase up to 100% (Figure 1).
- The frame losses on the traffic from stations towards APs have higher variability than those on the traffic from APs to stations (Figure 1).
- Some cards with Intersil firmware show *AP loss*. The card is set to a specific channel, but cannot see the frames from some APs on the same channel (Table II).
- The cards with Lucent firmware, in the monitoring mode, have *type loss* on ACK/RTS/CTS frames (Table III).
- Some sniffer cards and some APs show *incompatibility*, which leads to significant frame losses (Table IV).

We apply our passive measurement technique to infer the causes of end-to-end delays. We show two cases of abnormally high end-to-end wireless delays. Using our wireless measurement system we are able to identify the causes to be signal strength and card variability.

II. METHODOLOGY

We perform our experiments in the A.V. Williams building, at the University of Maryland (where the Department of Computer Science is located). The building has 58 access points installed which belong to three different networks, *umd*, *cswireless* and *nist*. *umd* network consists of 29 Cisco Aironet A-340 APs, which are set up by the Office of Information Technology of the University of Maryland. The *umd* network is most frequently used by the wireless users. *cswireless* (12 Lucent AP's) and *nist* (17 Prism2-based AP's) are built by individual research groups in the Computer Science department.

We set up several sniffer machines to capture wireless frames on the air. All sniffing devices used Linux operating systems with kernel version 2.4.19. We used *Ethereal* (version 0.9.6) and *libpcap* library (version 0.7) with the *orinoco_cs* driver (version 0.11b) as sniffing software. We made use of the 'monitor mode' of the card to capture the 802.11 header as well as physical layer header, called *Prism2* monitor header.

A. Captured Data

Prism monitor header is not a part of 802.11 frame header, but is generated by the firmware of the receiving card. The header includes RSSI(Received Signal Strength Indication), SQ (Signal Quality), Signal strength and Noise (in dBm) and Data rate (in Mbps).

IEEE 802.11 frame structure incorporates the following fields: protocol version, frame type (management, data and control), *Duration* for Network Allocation Vector (NAV) calculation, BSS Id, *Source* and *Destination* addresses, fragment, sequence number etc [4].

In this study we consider a frame to be "From-AP", if the frame is being transmitted by an AP to some wireless station. Sequence numbers in From-AP frames are generated by the source AP. Likewise "To-AP" frames are being transmitted by a wireless station to some AP. To-AP sequence numbers are generated by the source wireless station.

Card	From-AP Loss				To-AP Loss				Total
	# Distinct	# Retrans	# Loss	% Loss	# Distinct	# Retrans	# Loss	% Loss	
Linksys	26404	448	4098	13.44%	3176	269	66%	2.04%	12.34%
Lucent	20211	412	2614	11.45%	3055	271	121%	3.81%	10.52%

TABLE I

FRAME LOSS IS CALCULATED BY SEQUENCE NUMBERS. THE TRAFFIC USED IN CALCULATION ASSOCIATES WITH ONE AP.

III. PASSIVE OBSERVATIONS

A. Frame loss, AP loss and type loss

In order to estimate the frame losses we need to have an idea about the whole traffic. As we can not guarantee that at any point we have such complete information about the traffic, we use the approach of inferring the missed frames from the information we gather from the sequence numbers. Frame loss Number (*# Loss*) is the sum of all the gaps between any two adjacent sequence numbers.

In Table I, we calculate number of distinct frames (*# Distinct*), number of retransmissions (*# Retrans*) and frame loss number (*# Loss*) for two different wireless cards. As the maximum sequence number is 4095, we consider the wrap-around to calculate a gap. Frame loss rate (*% Loss*) is calculated by

$$\% Loss = \frac{\# Loss \times 100}{\# Distinct + \# Loss} \quad (1)$$

In Table I From-AP loss rate is based on AP-generated sequence numbers, while To-AP loss rate is calculated by summing up all the loss from all the wireless stations.

If the card cannot detect *some* packets from specific remote stations, this method of determining losses can lead to incorrect result. For example, suppose a wireless card can correctly detect 100 frames from node 1, without any loss, but cannot correctly detect *any* of the existent 100 frames from another node 2. The frame loss rate is calculated to be 0% incorrectly by considering only node 1 frames. However the correct loss is 50% (100 frames are lost out of 200).

In this study we use basically sequence number technique in loss calculation. Therefore when comparing two cards we confirm that the two cards have correctly detected the same set of wireless stations.

In Table II "*# Frames*" indicates the number of captured frames that associate with a specific AP. "*Pct.*" indicates its percentage among all the captured frames. Table II shows that Linksys card on channel 11 loses the frames of AP4 and AP5, which are on the same channel as the card (an example of *AP Loss*. Note that Linksys card also sees other channel AP (AP2 on 6).

In Table III, Lucent has *Type Loss*. Lucent loses the control frames of types ACK, RTS, CTS and Power-save. Actually, the receiving Lucent device captures the frames of such types but its firmware does not pass them up to sniffing software.

Both Table II and Table III are obtained from the same two day long experiment.

B. Loss variability

In Figure 1 frame loss rate varies between 0% and 100% during 4 day passive measurement experiment. In this plot we observe that the To-AP traffic has more loss variability.

In Table IV we group the traffic by the associated APs. Clearly frame losses vary depending on the card used in sniffing device as well as the AP which is being monitored. As before, we see significant variability between the traffic observed by the the two wireless cards.

<i>AP</i> (<i>essid,channel</i>)	<i>Linksys</i>		<i>Lucent</i>	
	<i># Frames</i>	<i>Pct.</i>	<i># Frames</i>	<i>Pct.</i>
AP1 (umd,11)	2583659	84.5%	2550568	41.3%
AP2 (nist,6)	454630	14.9%	6391	0.1%
AP3 (nist,11)	18579	0.6%	1172182	19.0%
AP4 (umd,11)	0	0.0%	1320012	21.4%
AP5 (nist,11)	11	0.0%	895638	14.5%
<i>Other rows are omitted</i>				
Total	3058516	100.0%	6182077	100.0%

TABLE II

NUMBER OF FRAMES ASSOCIATED WITH EACH AP: LINKSYS SNIFFER HAS AP LOSS ON AP4 AND AP5, WHICH ARE ON THE SAME CHANNEL AS THE LINKSYS.

<i>802.11</i> <i>Frame type</i>	<i>Linksys</i>		<i>Lucent</i>	
	<i># Frames</i>	<i>Pct.</i>	<i># Frames</i>	<i>Pct.</i>
Data	888082	25.9%	1318942	21.3%
Beacon	2117923	61.9%	4712323	76.2%
ACK	323674	9.5%	0	0.0%
RTS	34729	1.0%	0	0.0%
CTS	6734	0.2%	0	0.0%
Probe	52447	1.5%	150796	2.4%
Power-save	44	0.0%	0	0.0%

TABLE III

NUMBER OF FRAMES WITH EACH 802.11 TYPE: LUCENT SNIFFER HAS TYPE LOSS ON ACK/RTS/CTS/POWER-SAVE.




Fig. 1. Loss varies up to 100% over time: To-AP traffic is more variable than From-AP traffic. All the traffic associates with one AP.

AP(essid, Ch.)	Card	From-AP Loss				To-AP Loss			
		# Distinct	# Retrans	# Loss	% Loss	# Distinct	# Retrans	# Loss	% Loss
AP1(umd,6)	Linksys	4675	41	2	0.04%	210	9	16	7.08%
	Lucent	4656	41	17	0.36%	223	13	4	1.76%
AP2(nist,1)	Linksys	3085	17	78	2.47%	0	0	0	0.00%
	Lucent	1	0	0	0.00%	0	0	0	0.00%
AP3(cs,6)	Linksys	694	3	2414	77.67%	0	0	0	0.00%
	Lucent	2840	14	300	9.55%	0	0	0	0.00%
AP4(umd,6)	Linksys	4737	37	110	2.27%	249	9	114	31.40%
	Lucent	4701	38	144	2.97%	381	14	79	17.17%

TABLE IV

FRAME LOSS VARIES OVER THE CARD AND THE ASSOCIATED AP: ALL THE TRAFFICS ARE MEASURED IN THE SAME EXPERIMENT.

CARD VARIABILITY AFFECTS FRAME LOSS.

IV. APPLICATION: DIAGNOSIS ON END-TO-END DELAY

A. Active Experiment Setup

For diagnosis of wireless network delay, we conduct two-way UDP packet exchange experiments using NetDyn tool [2]. *Source* wireless station sends 20000 packets with the full payloads (1472 bytes for UDP) to a wired server (called *echo*). *Source* also acts as a *sink* machine to receive the packets back from *echo*. The *Source* node puts a sequence number and a time stamp in the payload. *Echo* node also puts a sequence number and a time stamp before forwarding the packet to the *Sink* node which puts a receive time stamp. Our sniffing devices capture the wireless frames between the source/sink wireless station and the access point.

Therefore, using NetDyn tool and the sniffers we have three NetDyn-generated timestamps (by *source*, *echo* and *sink*) and two sniffer-captured wireless traces (From-AP and To-AP) of all the successfully returning packets.

B. Delay Analysis

We present two high RTT (Roundtrip Time) delay cases where we place the source machine at different locations. We then analyze the two cases using our sniffer traces. Figure 2 and Figure 3 show RTT graphs of Case 1 and Case 2 respectively, which are measured by subtracting *source* timestamp from *sink* timestamp.

In Case 1, there exist high RTT delays up to 0.8 seconds and 57% packet loss. Case 2 is the case in which RTT delays are as high as 2.3 seconds and 75% of packets are lost. Note that in this section by *packet loss* we mean the loss experienced by the application in wireless station, while we mentioned the *frame loss* in the sniffers in the previous section.

To infer the causes of Case 1 delay we depict *source*, *echo* and *sink* timestamps and From-AP, To-AP timestamps in Figure 4. We can see that most delays occur between *source* and *echo* periodically. There are no such high delays on returning path, between *echo*, sniffers and *sink*. In the sniffer trace we find many retransmissions from *source* to the AP. We can conclude that weak signal strength causes many retransmissions, which leads to losses and delays in *source* machine's sending buffer.

In Case 2 the situation is more complex. From the timestamps and traces (as in Figure 5 and Figure 6) we see that most delays reside on wireless From-AP (returning) path. To infer the causes of the wireless delay we use physical layer information available in the sniffer trace. We find that From-AP data rate was consistently varied with the changing signal conditions as [8] proposes. If the AP does not receive ACK from the source at 11 Mbps, then the AP retries to send the packets at lower data rate (5.5 or 2 Mbps). But, due to the card incompatibility, the source card fails to receive the packets at the lower data rate, and incurs




Fig. 2. Case 1: High RTT delays up to 0.8 seconds and 57% packet loss.

retransmissions from the AP to *source*. Therefore delay becomes increases, becoming up to 2 seconds. As in Figure 7, the *source* card send/receive only at 11Mbps.

In Figure 8, we can see that there is high signal strength variability on wireless path. Such signal strength variability, combined with card incompatibility brings the high RTT delays in Case 2.

V. CONCLUSION AND FUTURE WORK

In this study we have extensively used passive measurements of the wireless medium to identify loss and loss variability in the sniffers as major pitfalls in wireless monitoring and measurements. We observed various kinds of losses, such as *frame loss*, *AP loss* and *type loss*, and their variability during our experiments.

We also presented some active experiment results to explain how signal strength and card variability cause frame loss and end-to-end delay between the AP and a wireless station. The wireless measurements used in this experiment effectively identify that the causes of such delay are signal strength and card variability. Through our experiments we have shown how careful wireless monitoring can be used to infer richer information about both the wireless medium and how the medium affects traffic patterns, which is not possible with a measurement from a wired vantage point.

Currently we are working on 802.11 conformance test with sniffer traces. 802.11 conformance test aims to detect misbehaviors or bugs in actual 802.11 implementations. With the traces collected by a single or multiple sniffers, we can obtain a detailed description of the misbehaviors of some nodes. In our preliminary experiments we have identified such implementation inaccuracies using our monitoring system.

REFERENCES

- [1] A. Balachandran, G.M. Voelker, P. Bahl, and Rangan P.V. Characterizing User Behavior and Network Performance in a Public Wireless LAN. In *Proceedings of SIGMETRICS*, June 2002.
- [2] S. Banerjee and A. Agrawala. Estimating Available Capacity of a Network Connection. In *Proceedings of IEEE International Conference on Networks*, September 2001.
- [3] B.J. Bennington and C.R. Bartel. Wireless Andrew: Experience building a high speed, campus-wide wireless data network. In *Proceedings of MOBICOM*, September 1997.
- [4] IEEE Computer Society LAN MAN Standards Committee. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. In *IEEE Std 802.11-1999*, 1999.




Fig. 3. Case 2: High RTT delays up to 2.3 seconds and 75% packet loss.




Fig. 4. Case 1: Source, echo, sink timestamps (by NetDyn), From-AP, To-AP timestamps (by sniffers). Delays exist between source and echo every 0.5 second periodically. No high delays exist on wireless path.




Fig. 5. Case 2: Source, echo, sink timestamps. Delays exist between echo and sink.




Fig. 6. Case 2: To-AP/From-AP traffic is captured by the sniffers. Delays reside on *wired* echo-AP path or *wireless* AP-sink path.

Time vs. Data rate at (60, 0) (021031_011754, zapp_prism2, gap=5ms, pkt size=1472, loss=75%)




Fig. 7. Case 2: AP varies data rates at 11, 5.5 and 2 Mbps (From-AP data rate, graph on top). Source cannot synchronize with the AP, send/receive packets only at 11Mbps (To-AP data rate, graph at bottom).

Time vs. Signal strength at (60, 0) (021031_011754, gap=5ms, pkt size=1472, loss=75%)




Fig. 8. Case 2: High variability in signal strength is observed by sniffers, which causes AP to shift data rate adaptively.

- [5] D. Eckardt and P. Steenkiste. Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network. In *Proceedings of SIGCOMM*, August 1996.
- [6] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *Proceedings of MOBICOM*, September 2002.
- [7] D. Tang and M. Baker. Analysis of a Local-Area Wireless Network. In *Proceedings of MOBICOM*, August 2000.
- [8] J.C. Wu, H. Liu, and Y. Lung. An Adaptive Multirate IEEE 802.11 Wireless LAN. In *Proceedings of the 15th International Conference on Information Networking*, January 2001.